Amplifier Anatomy - Part 1

By: Patrick Quilter, Chief Technical Officer
QSC Audio Products, Inc.

Article in Sound & Video Contractor Feb. 20, 1993

What do power amplifiers do? Power amplifiers drive loudspeakers. After an audio signal has been mixed, equalized and otherwise processed at a standardized line level, it is sent to the power amplifier. Its job is to increase the power of the signal until we get the desired sound level from the loudspeakers, without otherwise altering the waveform of the signal.

We need to talk about loudspeakers for a moment. A loudspeaker is an electromagnetic device that converts electric current into motion at audio frequencies. Because of the weight og the cone and the unavoidable resistive losses in the voice coil, it takes a lot of power to produce a high sound level. We know (especially if we read last month's article, "Mr. Ohm and His Talking Electrons") that electric power is a product of voltage and current.

Loudspeaker cone movement is proportional to the current in the voice coil. The amount of heat in the amplifier components is also proportional to current. However, it takes voltage to make the current flow, so a power amplifier must deliver high voltage and current simultaneously. Loudspeakers are generally made with a voice coil resistance of about 8 ohms, so the amplifier must produce 8V across the loudspeaker terminals to cause a 1A current to flow. This means that ideally the amplifier works into an impedance of 8 ohms.

In the real world, the impedance of loudspeakers is more complex. As the cone moves, it generates electrical back pressure, which can increase or decrease the current flow from the amplifier. Because of interactions with air pressure in the loudspeaker cabinet, the cone motion varies greatly, especially in the bass region. Therefore, the loudspeaker's impedance varies at different frequencies, and may range from 4 ohms to 20 ohms, averaging about 8 ohms. In addition, more than one 8 ohm loudspeaker may be connected to the amplifier. For these reasons, most professional amplifiers are built to work into impedances as low as 2 ohms, which will draw up to four times the normal current. The amplifier must also withstand very high impedances in case the load is disconnected. This requirement is normally not a problem because the current flow with no load is zero.

AMPLIFIER POWER

We all know that the amplifier's power rating tells us how loud the amplifier will get. A "200W at 8 ohm" amplifier is designed to deliver 40V to an 8 ohm loudspeaker, resulting in 5A of current (40V divided by 8 ohms). Of course, 40V times 5A yields the amplifier rating of 200W.

If we want to double the cone motion, we have to double the current, from 5A to 10A. Because the loudspeaker impedance is still 8 ohms, it will take 80V to get 10A. Therefore, the power rating must increase from 200W the 800W (80V multiplied by 10A). You can see why the power ratings escalate pretty quickly in high-output systems.

HOW DO POWER AMPLIFIERS WORK?

An amplifier is basically an ac-to-dc power converter. It takes ac power from the wall outlet (at fixed frequency and voltage) and converts it to audio power at the loudspeaker terminals (with variable frequency and voltage). The audio output is supposed to be a faithful replica of the line-level audio input, only larger.

Let's look a little further into the block diagram, at some of the major sub systems inside a power amplifier. We will explain more about each sub-system later in the article.

First we need a power supply. This subsystem accepts the ac power from the wall, isolates the audio circuitry from shock hazard, raises or lowers the ac voltage to suit the needs of the amplifier power rating, converts the ac power to dc, and stores it in an energy reservoir.

The other major subsystem is the output section. This is the electronic circuitry that accepts the line-level audio input and uses this information to control high-power transistors. These convert the energy contained in the dc reservoir to a high-power audio waveform that is a magnified replica of the input signal

AMPLIFIER PERFORMANCE LIMITATIONS

All amplifiers have a maximum power limit. The voltage at the amplifier output can only go as high as the voltage in the dc power supply. If the signal tries to exceed this limit, it "hits the ceiling," and the waveform becomes flattened. This problem, called clipping because it looks like the top of the waveform has been clipped off, results in the familiar "blatting" sound of an overdriven amplifier. Increasing the supply voltage adds cost and weight to the amplifier, so amplifier power has a big effect on price.

Amplifiers have a minimum rated output impedance, which should be equal or less than the impedance of the loudspeaker load. As the impedance of the loudspeaker gets lower, more current will be drawn from the amplifier. This is why, up to a point, the amplifier power rating increases into lower impedances. However, the increased current puts a greater strain on the amplifier components and the power supply. At some minimum impedance, the strain will get so high that the power-supply voltage sags or the transistors overheat. Any further decrease in impedance will cause the amplifier circuitry to collapse, resulting in less power, or it could even cause amplifier failure.

Amplifiers also must reproduce all audio frequencies, from the highest to the lowest, at equal volume. This ability is called flat frequency response because the graph of amplifier gain vs. frequency is a flat line.

If the gain at low frequencies falls off, the sound will be thin or lacking in impact. If the high frequency gain rolls off, the sound will be dull or muffled. Most modern direct-coupled amplifiers are capable of very flat response, but sometimes the frequency response is intentionally limited to protect the loudspeakers from excessive power at frequencies we can't hear.

MORE ABOUT THE POWER SUPPLY

Why do we have to convert the ac power from the wall into dc power, and then back to ac? The ac power from the wall is at a fixed voltage and frequency, which are completely different from the audio voltages and frequencies. If we tried to use the ac voltage "as is" for a power supply, we would only be able to reproduce small parts of the audio waveform.

We have to convert the ac power into a fixed dc source, and provide enough energy storage to carry us through the periods where the ac voltage is passing through zero. This way, the audio output section has the power available to respond at anytime as required by the input signal.

So, let's take a more detailed look at the amplifier's power supply.

The ac power comes into the amplifier through the ac cord, is controlled by the on/off switch, and usually goes through a fuse or circuit breaker, which cuts off ac power in case of massive overload. It then reaches the power transformer, which is in the heart of the power supply. A transformer consists of two coils of wire around a common magnetic core.

The ac power is connected to the first or primary winding, which converts electric energy into magnetic energy. This magnetism flows through the iron core to the secondary coil, which converts it back to electricity.

Why do we do all of this? The two coils are insulated from each other, so that the secondary coil is isolated from any shock hazard in the primary coil. The ac voltage and current can also be scaled up or down by changing the number of turns in the secondary coil. Transformers are son useful that they are the major reason we use ac power distribution instead of dc (transformers only work on ac).

The simplest and least expensive transformer is the E-I type, which is generally cubic-shaped (roughly equal height, length, and width). This type is widely used, but it has a tendency to give off hum, which might be picked up by nearby circuitry.

The U-I type is more expensive, but it is easier to make in a flatter shape that can fit into low-profile amplifiers. It also reduces the hum emissions. The toroidal type is built on a donut-shaped core, which has the best magnetic properties. It can be made quite flat, it weighs somewhat less and is has low hum emissions, but it is the most expensive. At QSC, we use the U-I transformer for most of our low-profile amplifiers because it offers most of the advantages of the toroid at a lower cost.

Once we have scaled and isolated the ac power through a transformer, we need to convert it to dc. This is the job of the rectifier and dc filter capacitors.

The rectifier is a one-way ratchet that takes the back-and-forth flow of ac current and redirects it so that is always flows in the same direction. The rectifier uses diodes, which permit current flow in one direction and block flow in the reverse direction.

A full-wave bridge rectifier circuit uses four diodes. Each diode passes current only in the direction of the arrow. If you follow the current flow around the circuit, you will see that, no matter which way the ac current is flowing into the rectifier, it always emerges in the same direction. Now we have the dc instead of ac (the current only flows in one direction), but it still has big valleys in it. We need to smooth out this ripple voltage. This is the job of the filter capacitors.

Capacitors are like tanks that hold electricity. Once you fill them, it takes a while to drain them out. Therefore, we connect a large capacitor to the output of the rectifier. The capacitor fills, or charges up, to the peak voltage of the rectified wave-form. If the capacitor is large enough, it stays pretty full between the peaks, and we get an almost perfectly smooth dc voltage. The electronic value of the filter capacitors determines how well the ripple voltage is removed.

In the last 20 years, high-value capacitors have been considerably reduced in size. The high-density filter capacitors are easier to mount on the circuit board right next to the power transistors, which helps improve the high-frequency performance of the amplifier. The length of wiring between the old-style capacitors and the transistors can introduce a slight inductance or electronic lag, which prevents the transistors from instantly drawing power from the supply.

POWER-**SUPPLY REGULATION**

Another concern with power supplies is regulation: the ability to hold the dc supply voltage constant, despite changes in amplifier loading or ac voltage.

The first characteristic is called load regulation, or sometimes power-supply stiffness, and basically depends on the resistance in the transformer. An ideal transformer would have zero resistance and would be able to maintain a constant voltage (perfect regulation) no matter how little or how much current the amplifier needs. Real-world transformers have resistance in the wire coils, which causes the supply voltage to drop when current flow increases. To minimize this voltage drop, thicker wire must be used, which increases the size and weight of the transformer.

Earlier solid-state amplifiers used rather large transformers to keep the no-load voltage (at rated power). The designers needed to minimize no-load voltage because high-voltage transistors were expensive, if not impossible to get. In recent years, the cost of high-voltage transistors has come down, so the trend has been toward somewhat smaller transformers to reduce the weight of amplifiers, even though the no-load voltage rebounds to a higher level.

A side effect of the voltage drop is that, because the filter capacitors charge up to higher voltage during periods of low demand, the amplifiers can deliver a momentary burst of power above its normal rating. This feature, called dynamic headroom, can add 2dB or 3dB of peak undistorted power, which is equivalent to having up to 100% more wattage.

The ability to hold a constant voltage despite ac voltage fluctuation is called line regulation. Ordinary passive power supplies, such as the transfomer-rectifier-capacitor system discussed earlier, do not offer line regulation. The dc supply voltage changes along with any change in the ac voltage. The power company usually tries to maintain a constant voltage, but heavy loads or long ac cables can cause voltage drops, which result in loss of amplifier power.

Companies that use a great deal of amplification, such as touring companies, must invest a lot of money in heavy-gauge ac cabling (distro systems) to minimize this effect. Until switching power supplies become more practical, the correction to this problem would unfortunately add cost and weight to the amplifier, so the tendency has been to spend the same money making bigger amplifiers. You get the desired minimum power under worst-case conditions, and you come out ahead when ac service is normal.

SWITCHING POWER SUPPLIES

The size and weight of power-supply components has been somewhat reduced over the last 20 years, but progress has been slow because we are only refining the same basic technology. Meanwhile, other industries, such as the computer industry, have been perfecting light-weight switching supplies reduce the size and weight of the power transformer by operating it at a much higher frequency. For reasons beyond the scope of this article, high-frequency transformers are much smaller that low-frequency transformers. However, we are stuck with the 50Hz or 60Hz ac power supplied by the power company, so if we want to use high-frequency transformers, we must generate our own high-frequency power, which results in a fairly complicated block diagram.

First, we rectify the incoming ac and smooth it with capacitors, just as we did with the passive supply described above, but without an ac transformer. Then we use a high-speed switching transistors to convert the dc power to a high-frequency ac waveform, usually 50kHz to 100kHz (about 1,000 times higher than normal ac power). This high-frequency ac is fed to a small high-frequency transformer, which isolates the secondary from ac shock hazard and scales the voltages, just as the large ac transformer did in the passive supply. This high-frequency ac voltage is the rectifier and filtered again, resulting in the final dc supply for the amplifier. The active supply is much more complicated than the passive supply, but the weight of the components is much less.

Although active supplies are more expensive, costs are slowly coming down, and there are important advantages. In addition to the primary benefit of greatly reduced weight, we can control the operation of the high-frequency transistors to compensate for variations in ac voltage and load currents, thus improving both kinds of power-supply regulation. The ultimate result will be more consistent amplifier performance, but the audio industry must solve problems of cost, reliability and radio/TV interference caused by the high-frequency switching. This will undoubtedly be an active area of progress in the decade of the 1990s.

MORE ABOUT THE AUDIO OUTPUT CIRCUIT

The story of the actual power amplifier circuit begins with the input connectors. These humble components are crucial to getting a high-quality signal into the amplifier because garbage in is garbage out. In addition to corrosion-proof plating and strong mounting, it’s a big help to have balanced inputs, which are now fairly standard. Balanced inputs permit the amplifier to ignore most forms of interference that occur in the cabling between electronic units.

Most amplifiers also have a gain control. It is usually operated full up, but it is handy to be able to reduce gain for testing or to lower the noise floor when you know you have input volume to spare. After the balanced-input and gain-control circuitry, we enter the actual power amplifier circuit.

The main function of this circuit section is to increase the input signal from about 1V to about 100V, and to increase the current from about 0.1mA to about 30A. This is a power gain of about 30 million! To understand how this occurs, we need to discuss how transistors work.

Transistors (and tubes in earlier years) are variable-resistance elements that are connected between the dc supply and the load (the loud-speaker). The transistor acts as a valve. A small input signal causes a much larger amount of current to flow from the dc supply to the load. The device controls a load current about 50 to 100 times greater that the input current, so the device has a gain of 50 to 100.

To increase the gain, we can cascade devices by driving a second transistor with the output of the first transistor, and so on. This way, we can build up the tremendous gains we need. The exact method of cascading is on of the major differences among amplifier designs, and it would double the length of this article to fully review these methods, However, we can give some basic terminology.

The last, or highest-power, set of transistors are called output transistors. These are high-power devices mounted on large heat sinks. The outputs are driven at a much lower power by driver transistors. Sometimes there are pre-drivers before the drivers, though in some cases it is possible to go to small-signal devices. QSC amplifiers use a very-high-gain integrated circuit (called an op-amp), followed by a relatively simple 2-stage set of driver and output transistors.

The large load current is ideally a magnified replica of the small input current. However, for a number of reasons, the load current might no be an exact replica; it might be distorted. The most obvious kind of distortion is clipping, which occurs when the voltage across the load comes so close to the dc voltage that the transistor saturates, or bottoms out, and can't go any further. A lesser form of distortion occurs because the transistor's gain is not uniform: It varies because of temperature and current differences.

All of these effects are called non-linearities because the transistor deviates from the ideal of uniform magnification, much like wavy glass makes straight lines look crooked. We will explain how distortion is minimized later in the article.

Another major problem with transistors is that they are one-way devices: They only handle positive or negative currents. Therefore, we need a way to connect positive and negative devices together to deliver a complete audio waveform. This method is called push-pull operation and has been the key to high power performance since early tube amplifiers. There are a number of ways to combine push-pull currents.

Yet another problem is that of heat loss. Let's say we connect a transistor to a 100V supply, but for the moment we only ask it to deliver 40V to an 8 ohm load. We know (from the example used earlier) that 5A of current will flow in the load, resulting in 200W of output power. That same 5A must flow through the transistor to get to the load. At the same time, the “unused” 60V appears across the transistor. We have a combination of 5A and 60V in the transistor, which results in 300W (5A x 60V) of power in the transistor.

A basic low of physics states that energy cannot be destroyed, only changed. Because we aren't letting this unused power go to the load, it has to go somewhere, and the result is waste heat. This example show that it is very easy for the power wasted in transistors to exceed the power delivered to the load. This waste heat is the reason powerful amplifiers need large heat sinks, which dissipate the unwanted heat. Otherwise, the transistors would get so hot that they would fail.

Although I won't burden the reader with excess detail about different ways to cascade transistors in amplifiers, the final push-pull output transistors can be combined in general ways that affect distortion and heat loss. These categories, called classes of operation, were defined many years ago to allow discussion of these trade-offs. You have probably heard of class A, class B, class AB and other lettered designations for amplifiers. Here's a brief explanation of their meaning (and it has nothing to do with USDA meat grades).

CLASS A

This is the easiest class to understand, so it leads the list. The positive and negative output transistors each handle 100% of the audio signal- they are biased so their zero-signal output current idles halfway between zero and maximum. When the audio current in one transistor increases, the current in one transistor increases, the current in the other decreases; as a result, their voltage move together. Each transistor can, therefore, deliver a faithful replica of the signal all by itself, except for the large idle current.

If we were to connect just one transistor to the loudspeaker, we would hear fair-quality sound, but that would force the cone way off center and would probably overheat the voice coil. When we connect both transistors to the load, the idle current from one transistor is absorbed by the other (rather than going through the load), but the audio currents reinforce each other and appear nicely in the load.

The primary advantage of class-A operation is inherent lack of distortion. The full waveform is preserved in the positive and negative transistors, so there is no trick to combining their currents.

However, a serious flaw is the extreme heat loss at idle. The transistors actually run hottest while “standing still”- sort of like controlling a car’s speed with the brakes while keeping the throttle mashed down. Naturally, amplifier designers have looked for other ways to ensure low distortion without such wasteful operation.

CLASS B

If we are careful, we can let each transistor control only its half of the waveform. When the waveforms are combined properly, we still get the complete output waveform, but we have eliminated the large idle current. The amplifier runs much cooler because no power is used until it’s needed.

The trick, of course, is to get seamless combining. If the waveforms don’t joined together perfectly, we get zero-crossing distortion (frequently called crossover distortion). This kind of distortion is quite objectionable because it results in a slight gargling or rattling sound during quiet parts of the program, where the signal is near zero.

Fortunately, there are a number of ways to eliminate this problem. One popular method is to compromise between class A and B and operate the amplifier in class AB. Bu permitting a small idle current to flow, we get a small amount of idle heat, but we eliminate any chance of “dead space” between the positive and the negative waveforms.

CLASS C

When each transistor controls less than 50% of the waveform, we call this mode class C. This mode is not usable for audio because of the large gap between the waveforms, which causes severe zero-crossing distortion. Class C is used where such distortion is unimportant or can be tuned out by other circuitry. In some amplifiers, the output transistors are run in class C for less idle heat, with the driver transistors filling in the gap. This method is called class ABC.

CLASSES D, E, & F

These classes apply to switching amplifiers, which will be explained in the second half of this article.

CLASS G

This mode uses two or more sets of output transistors connected to different supply voltages. The goal is to reduce the heat loss in class A or B amplifiers. Remember the example in which we had a 100V power supply but we only needed 40V into the load? We had a lot of waste heat because there was 60 “unused” volts wasted in the output transistors.

In a class G amplifier, we have on set of transistors connected to a lower voltage supply, say 60V, which supplies all output voltages up to this value. We then transfer to a second set of transistors connected to the 100V supply. With this method, the ”unused” voltage for a 40V output is cut from 60V to 20V, dramatically reducing the waste heat. Because an amplifier spends most of its time supplying only a fraction of its power, the average losses can be cut by 50% or more.

The main problem is to ensure seamless transfer from the low-voltage to the high-voltage transistors to avoid any small glitches similar to zero-crossing distortion. QSC Series Three and the original MX series amplifiers used this technique quite successfully.

CLASS H

This class uses a single bank or output transistors connected to a low-voltage supply, along with some means of switching them to a higher-voltage supply when required. This method has the same thermal benefits as class G, but it avoids the second bank of output transistors, thus reducing the size and cost of the amplifier.

The QSC EX series uses this technique to pack more power in the same chassis. (The EX4000 has twice the power of the old MX2000.) The new MXa series uses the same technique to simplify the construction and to improve reliability.

Most of these methods require some alteration of the audio signal as it is broken apart and reassembled. Not surprisingly, these alterations result in errors in the reassembled waveform. In Part II, we’ll discuss how error-correction circuitry, protection circuitry and other design features allow the amplifier to perform its function without altering the waveform.

Amplifier Anatomy - - Part 2

Article in Sound & Video Contractor, March 20, 1993

Power amplifiers, those “unsung heroes” of the sound system, have traditionally been the old reliable that consultants and contractors learned to count on as everything else in the sound system became more complex. While the rest of the system demanded more and more attention, the power amps were something you just didn’t have to worry about.

Although power amplifier technology has more to offer and requires more thought, there are some basic principles that every leading amplifier manufacturer follows. Understanding how equipment designers are solving amplifier problems can give you a new appreciation for this basic piece of equipment.

In Part I, we described how an amplifier works, how power supplies affect amplifier function and which amplifier classes work best for various situations. In Part II, we will go into more detail about amplifier circuitry and design.

NEGATIVE FEEDBACK AND DISTORTION

In Part I, we mentioned that transistors are not inherentlyperfect magnifiers. Most of the advanced circuit techniques we described involve re-assembly of the audio waveform in various ways. If we had no way of correcting errors, the result would typically be a rather garbled and harsh reproduction. Fortunately, we have a powerful error-reduction technique called negative feedback.

You probably associate negative feedback with criticism. Actually, this perception is not so far-fetched, but in electronics there is nothing bad about “negative” feedback. This technique is basically the same process we use daily when observing one’s progress and making mid-course corrections.

Let’s use driving as an example. Imagine if you tried to steer around a corner with your eyes closed. Even if you turned the wheel about the right amount at about the right place before you closed your eyes, the car would soon leave the road because of uncorrected small errors. In the real world, you drive with your eyes open. You turn the wheel, observer how the car is turning, and then make small corrections to maintain the desired course. This is a basic example of how we use feedback to correct for errors and produce the desired result.

We can also use feedback in amplifiers. The output of real-world circuitry is distorted: The output is higher or lower than desired. We usually don’t know just what the errors are, or we could correct for them. We can correct for unknown errors by comparing the output to the input and telling the circuit to increase or decrease the output until they match.

To describe the actual process, let’s assume we want an amplifier with a gain of 10. The actual, imperfect amplifier has a gain that varies unpredictably from 8 to 12, resulting in up to 20% of output error. We attenuate, or reduce, the actual output of the amplifier by a factor of 10. If the amplifier were error-free, this reduced feedback signal would be a perfect match to the input signal, but the actual feedback signal varies around the input signal by 20%.

We have an accurate picture of just the error because we have compensated for the desired gain by the 10-to-1 attenuation. Now, here’s the tricky part. If we magnify the apparent error by amplifying the mis-match between the input signal and the feedback signal, then combine this information with the original input signal, we can get the main amplifier to reduce its own errors automatically.

We have a crucial choice, however. The error signal can be added or subtracted from the input. If we add the error signal, it just makes the error worse. This is called positive feedback, and it turns transistors into oscillators. It turns sound systems into oscillators, also. In an amplifier it leads to wild, runaway operation. If, on the other hand, we subtract the error signal, the error is always diminished. Thus, we have used negative feedback.

In practice, we reduce error by putting much more gain than we need into the amplifier. It we increase the gain of the amplifier by a factor of 10, the gain, even with errors, will range from 80 to 120. When we “close the feedback loop,” using the 10-to-1 attenuation, the error signals always a large positive value. The amplifier quickly reduces its own output until the feedback and input signals match up.

In effect, because the amplifier has extra gain, it is in a constant state of “holding back,” which makes it easier to hit the desired target. With enough extra gain, it is ultimately the accuracy of the feedback circuit itself, not the amplifier, that determines the accuracy of the final output.

There is a limit, however. All circuitry has a slight lag between input and output. If you increase the gain of the circuit too much, it will become too sensitive, and the combination of lag and feedback will cause hunting or oscillation around the desired value. This problem, called instability, limits the amount of feedback that can be used.

The only fundamental cure is to reduce the circuit lag by using high-speed components. There has been a lot of progress in this area in the last 20 years, and today's transistors are about 10 times faster. This progress probably explains how solid-state amplifiers have gradually eliminated the harshness that some listeners heard in early amplifiers, which used relatively slow power transistors.

Feedback can be applied around all of the cascaded elements in the amplifier. This process, called global feedback, is popular because it corrects all internal errors in one swoop. Some designers prefer a process in which they sub-divide the amplifier into several cascaded feedback loops, called local feedback. Certain forms of instability are easier to eliminate with the local feedback technique, but internal signal levels must be higher when one feedback loop feeds the next. QSC amplifiers use global feedback because it is less expensive and is easier to assure accuracy with only one set of critical feedback components to worry about.

One notable effect of high feedback is on amplifier clipping characteristics. Without feedback, transistors approach saturation gradually, giving the effect of cushioning the impact. This type of clipping is called soft clipping. With feedback, the output signal is forced to stay on track until the last possible moment, resulting in an abrupt impact, called hard clipping, which sounds more fuzzy than soft clipping. Some amplifiers feature low-feedback designs to smooth out the clipping.

Because it is so hard to reduce distortion without high feedback, the trade-off is often between clean amplifiers with harder clipping and slightly mushier amplifiers with smooth clipping. The choice depends on personal preference and on how much the amplifier will be overdriven.

In any case, it is important to avoid sticky clipping. Poor feedback circuit design can make the amplifier track its input too far, and then snap back and ring without damping. The amplifier should enter and leave clipping cleanly with no snapping or chattering.

PROTECTION CIRCUITRY

The lower the impedance of the load, the greater the current drawn from the amplifier, and the greater the heat generated in the output transistors. If too many loudspeakers are connected to the amplifier, or if the ends of the loudspeaker wire touch together by accident, the load impedance goes very low, and the current flow becomes dangerously high. If the flow is not limited, the output transistors will burn out. Therefore, amplifiers need some kind of short-circuit protection. There are many ways to protect against short circuit, but the trick is that you can't prematurely limit performance into normal loads. QSC amplifiers continuously monitor the load impedance. Loads above 2 ohms draw safe currents, and only normal voltage clipping occurs. Below 2 ohms, the maximum amplifier current is reduced if it exceeds the safe limit for more than a fraction of a second. This way, short peaks are permitted even into marginal loads, yet the amplifier is still protected against gross, sustained overloads.

Other common protective circuits include turn-on and turn-off muting, shut-down or muting in case of excessive temperature, protection against radio pickup (RFI), and dc fault protection, which shuts down the amplifier is a transistor loses control. The design of these circuits is just as important as the actual audio circuitry because they can make the difference between surviving an accident or winding up with burnt rubble.

SWITCHING AMPLIFIERS

Remember those classes D, E and F that we promised we would talk about in Part I? As you have seen, heat in the output transistors is a major problem, and it is inherent in the way linear amplifiers operate. Whenever the amplifier is delivering only part of its power to the load, waste heat is created.

There is another way of converting dc power into audio power that reduces the inherent heat losses. Because the losses occur when the output transistors are partially on, we avoid this state. We turn them fully on and send all of the dc power to the load, or we leave them fully off so that no power flows. In both cases, little or no power is wasted in the transistor.

To get the desired average power in the load, we rapidly switch the transistors on and off for the desired percentage of the time, with the time the transistors are on varying from 0% to 100%. The switching must occur much faster than the highest audio frequency if the averaging is to work correctly. This is the switching amplifier or class-D amplifier. (Classes E and F are special cases, as is class C, that apply mainly to non-audio uses.)

How does on-off switching drive a loud-speaker? The magnetic field in the voice coil does not collapse instantly when the amplifier switches off. The loudspeaker continues interpolating a fraction of the waveform while the amplifier is switching.

Class-D amplifiers are only now becoming practical. The speed requirements for the switching transistors are 50 to 100 times greater than for linear audio amplifiers. The high-frequency switching causes radio interference, and many practical problems must be solved to attain the same audio fidelity that we expect with linear amplifiers. When the switching causes radio interference, and many practical problems must be solved to attain the same audio fidelity that we expect with linear amplifiers. When the switching is perfected, we can expect the heat sinks to be about one-fourth the size, reducing amplifier size and cooling requirements. When combined with the switching power supply, we will have a size and weight breakthrough comparable to the transition from tubes to solid-state. It will take a great deal of experience to overcome the reliability problems associated with complex new circuitry. This, too, should be an active area for development in the 1990's.

MECHANICAL DESIGN

We have treated amplifier design as an electronic problem, but the physical size, location and ruggedness of the parts is at least as important. You can inspect these features by eye at trade shows or on other occasions when the insides are on display.

The power transformer is the single heaviest element and must be mounted securely. Also, look for adequate clearance around it to allow for cooling and shifting in case the amplifier is dropped.

The power transformer is the single heaviest element and must be mounted securely. Also, look for adequate clearance around it to allow for cooling and shifting in case the amplifier is dropped.

The height of the chassis is another important variable. Low-profile chassis allow you to mount a lot more amplifier power in a given amount of rack space, but such a design pots much more strain on the rack ears. The rack ears should be well-connected to the chassis in such a way that overstress does not damage or loosen some other critical element, such as the faceplate. Rear support is strongly recommended. Standard-height chassis provide a greater mounting surface and allow for more internal space around components, such as the transformer.

Fan cooling creates noise in the chassis, but it dramatically reduces the size of the heatsink. Because of the increased power ratings of modern amplifiers, it is harder to find convection-cooled units. Class-G and -H amplifiers reduce heatsink requirements (as will class-D designs eventually) and may help bring back convection-cooled amplifiers. Meanwhile, high-quality, variable-speed fans can minimize the cooling noise. You should always check fan noise carefully if people will be sitting near the amplifier.

There are a number of details you should check. For example, external connectors, controls and displays should be recessed to protect them from external damage, and lockout covers are sometimes available for extra security. All connectors should be high-quality. Make sure input and output connectors are firmly mounted and will resist the strain of somebody tripping over the cabling. Internal connectors also need to be rugged.

Sub-assemblies and circuit boards might shift as the amplifier bounces around, so excessive rigidity can be a problem. Card-edge connectors are frequently troublesome. Cable-type or shock-mounted connectors with some “give” are preferred.

Gold plating is frequently preferred on input connectors, but it might wear through after a few insertions. The quality of the underlying plating is probably more important.

Internal connectors should also have corrosion-proof plating. Gold plating is best for small signal connections, but it can burn through at high currents. Heavy-duty connectors should use precious-metal alloys or conductive oxide platings in which current flow actually improves the integrity of the contact.

You’ll find these principles used in QSC amplifiers and in amplifiers from other leading audio companies. Various engineers have their own opinions on how each problem is best solved, but, in many cases, it’s basically a question of continuously refining a chosen approach. The basic principles here should help users appreciate what goes into modern amplifiers, often the unsung heroes of sound systems.